

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellants:	Roger HANSEN et al.	§	Confirmation No.:	5369
Serial No.:	10/737,374	§	Group Art Unit:	2114
Filed:	December 16, 2003	§	Examiner:	Loan Truong
For:	Persistent Memory Device For Backup Process Checkpoint States	§	Docket No.:	200312027-1

SECOND APPEAL BRIEF

Mail Stop Appeal Brief – Patents
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

Date: December 6, 2010

Sir:

Appellants submitted a Notice of Appeal and Appeal Brief in this matter on June 30, 2010. The Office action issued on September 20, 2010 was a *sua sponte* reopening of prosecution in view of the Appeal Brief. Appellants hereby exercise their right to continue the appeal process by submission of this Second Appeal Brief. A Notice of Appeal is being filed concurrently herewith.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

TABLE OF CONTENTS

I.	REAL PARTY IN INTEREST	3
II.	RELATED APPEALS AND INTERFERENCES	4
III.	STATUS OF THE CLAIMS	5
IV.	STATUS OF THE AMENDMENTS.....	6
V.	SUMMARY OF THE CLAIMED SUBJECT MATTER.....	7
VI.	GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL.....	9
VII.	ARGUMENT.....	10
	A. Section 103 Rejections of Claim 1-8 and 39-41 Over Vis, Traverstat and Kano.....	10
	1. Claims 1, 3-8 and 39-41.....	10
	2. Claim 2	17
	B. Section 103 Rejection of Claim 9 and 38 Over Vis, Traverstat, Kano and DeKoning.....	20
	C. Conclusion	20
VIII.	CLAIMS APPENDIX.....	21
IX.	EVIDENCE APPENDIX	24
X.	RELATED PROCEEDINGS APPENDIX	25

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

I. REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 11445 Compaq Center Drive West, Houston, Texas, 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC. The Assignment from the inventors to HPDC was recorded on December 12, 2003, at Reel/Frame 014817/0808.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals or interferences.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

III. STATUS OF THE CLAIMS

Originally filed claims: 1-37.
Claim cancellations: 10-37.
Added claims: 38-41.
Presently pending claims: 1-9 and 38-41.
Presently appealed claims: 1-9 and 38-41.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

IV. STATUS OF THE AMENDMENTS

No claims were amended after the Office action dated September 20, 2010.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

This section provides a concise explanation of the subject matter defined in each of the independent claims, referring to the specification by page and line number or to the drawings by reference characters as required by 37 C.F.R. § 41.37(c)(1)(v). Each element of the claims is identified with a corresponding reference to the specification or drawings where applicable. The specification references are made to the application as filed by Appellants. Note that the citation to passages in the specification or drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element. Also note that these specific references are not exclusive; there may be additional support for the subject matter elsewhere in the specification and drawings.

The various embodiments are directed to a persistent memory device for backup process checkpoint states.¹ At least some of the illustrative embodiments are systems as in claim 1:²

1. A system for storing checkpoint data comprising:
a network interface to an external network; **{6, [0022], lines 1-11; Figure 1A, element 114}** and
a persistent memory unit coupled to the network interface,
{6, [0022], lines 1-11; Figure 1A, element 102}
wherein:
the persistent memory unit is configured to
receive the checkpoint data into a region
of the persistent memory unit via a
remote direct memory write command
from a primary process through the
network interface **{6, [0024], lines 1-8}**,
and to provide access to the checkpoint
data in the region via a remote direct
memory read command from a backup
process through the network interface
{6, [0024], lines 1-8}, wherein the
remote direct memory write command is
preceded by a create request for the

¹ Specification Title.

² Citations to the specification from this point take the form **{[page], [[paragraph]], lines [lines within the paragraph]}**.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

region {13, [0047], lines 1-15} {14, [0047], 15-16} and the read command is preceded by an open request for the region {13, [0047], lines 1-15} {14, [0047], 16-17}; and

the backup process provides recovery capability in the event of a failure of the primary process. {6, [0024], lines 4-6}.

Other illustrative embodiments are systems as in claim 2, having all the limitations of claim 1 and further reciting:

2. The system of Claim 1, further comprising:
a persistent memory manager configured to program the network interface with information used by the network interface to perform virtual-to-physical address translation. {12, [0044], lines 1-11} {13, [0047], lines 1-15}.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Whether claims 1-8 and 39-41 are obvious under 35 USC § 103(a) over Viswanatham et al. (U.S. Pat. No. 7,165,186, hereafter “Vis”), Traverstat et al. (U.S. Pat. No. 6,941,410, hereafter “Traverstat”), and Kano et al. (U.S. Pat. No. 7,222,194, hereafter Kano).

Whether claims 9 and 38 are obvious under 35 USC § 103(a) over Vis, Traverstat, Kano, and DeKoning (U.S. Pat. No. 6,691,245).

VII. ARGUMENT

A. Section 103 Rejections of Claim 1-8 and 39-41 Over Vis, Traverstat and Kano

1. Claims 1, 3-8 and 39-41

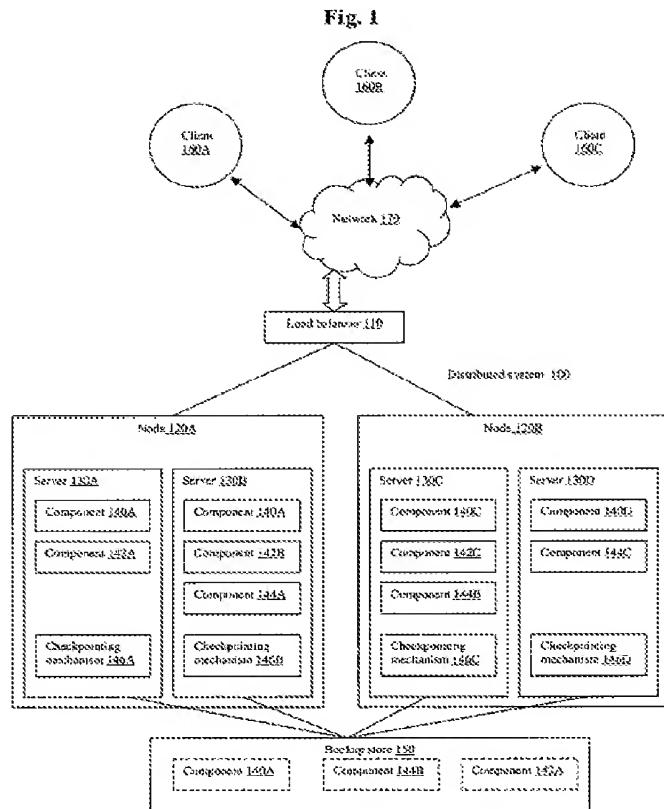
Claims 1, 3-8 and 39-41 stand rejected as allegedly obvious over DeKoning and Boyd. Claim 1 is representative of this grouping of claims. The grouping should not be construed to mean the patentability of any of the claims may be determined in later actions (e.g., actions before a court) based on the groupings. Rather, the presumption of 35 USC § 282 shall apply to each of these claims individually.

Vis is directed to selective checkpointing mechanisms for application components.³ In particular, Vis appears to be directed to a heuristic mechanism used to determine whether and/or when to perform checkpointing.⁴ With respect to the location to which checkpointing data is written, Vis appears to discuss only the backup store 150, shown at the bottom of Figure 1 of Vis reproduced immediately below.

³ Vis Title.

⁴ Vis. Col. 4, lines 17-28.

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010



More specifically, Vis describes a plurality of clients 160 (top of the figure) that couple to the distributed system 100 through the network 170 and load balancer 110.⁵ Separate and apart from the connections of the clients through the load balancer 110, Vis discusses the backup store 150 (bottom of the figure).⁶ In fact, in one case the backup store 150 resides within one of the nodes 120.⁷

Given the location of the backup store 150 in relation to the network 170 and load balancer 110, it seems clear that commands used to write checkpoint data to the backup store 150 do not flow through the network 170 and/or load balancer 110 of Vis. Vis appears to be silent as to the precise connection mechanism between the backup store 150 and the nodes 120.

⁵ Vis Col. 2, line 56 through Col. 3., line 10.

⁶ Vis Col. 3, lines 38-42.

⁷ Vis. Col. 3, lines 51-56.

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010

Representative claim 1, by contrast, specifically recites:

1. A system for storing checkpoint data comprising:
a network interface to an external network; and
a persistent memory unit coupled to the network interface,
wherein:

the persistent memory unit is configured to receive the checkpoint data into a region of the persistent memory unit via a remote direct memory write command from a primary process through the network interface, and to provide access to the checkpoint data in the region via a remote direct memory read command from a backup process through the network interface, wherein the remote direct memory write command is preceded by a create request for the region and the read command is preceded by an open request for the region; and

the backup process provides recovery capability in the event of a failure of the primary process.

Appellants respectfully submit that Vis, Traverstat and Kano fail to teach or suggest such a system.

**THE NETWORK 170 OF VIS IS NOT THE
PATH THROUGH WHICH CHECKPOINT COMMANDS FLOW**

In rejecting representative claim 1, the Office action makes the following statements:

In regard to claim1, Viswanatham et al. teaches a system for storing checkpoint data comprising:

*a persistent memory unit (backup store may be located on a separate computer from the nodes 120 A-B, fig. 1, col. 3 lines 53-56) coupled to the network interface (network interface may be a LAN, WAN, the internet or other types, col. 2 lines 63-67)...*⁸

It is clear from the rejection that the Office action relies on the backup store 150, but then relies on a network interface of Vis; however, no network interface is

⁸ Office action of September 20, 2010, Page 3, second and third full paragraphs (bold emphasis added, italics original).

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

specifically shown in Vis, and the Office action must therefore rely on an allegedly inherent network interface. The question then becomes, where would such an allegedly inherent network interface reside?

In answer to the question, the rejection goes further to rely on Col. 2, lines 63-67 of Vis, which cited location refers to the network 170 between the client devices 160 and the load balancer 110. Thus, the allegedly inherent network interface must reside between the load balancer 110 and the nodes 120. However, Vis clearly shows in Figure 1 the backup store 150 having a separate connection to the nodes 120 (bottom of the figure) than the load balancer 110 and network 170 (top of the figure). Thus, in order to arrive at the rejection the Office action must rely on an inherent teaching of Vis that in embodiments where the backup store 150 is “located on a separate computer from the nodes 120A-B” the connection to the separate connection would be through the load balancer 110 and network 170.

INHERENCY FAILS IN THIS SITUATION

The Manual of Patent Examining Procedures (MPEP) provides the following guidance as to inherency.

To establish inherency, the extrinsic evidence must make clear that the missing descriptive matter is **necessarily present** in the thing described in the reference, and that it would be recognized by persons of ordinary skill. Inherency, however, may not be established by probabilities or possibilities. **The mere fact that a certain thing may result from a given set of circumstances is not sufficient.**⁹

In situations of Vis where the backup store 150 is “located on a separate computer from the nodes 120A-B,” it simply is not necessarily present that the commands to write information to the backup store 150 would flow through the load balancer 110 and network 170. In fact, Figure 1 of Vis expressly teaches a configuration where the path across which data is written to the backup store 150 is different than the load balancer and network 170.

⁹ MPEP 8th Ed., Rev. 6, September 2007, § 2112(IV), p. 2100-47 (internal quotations omitted, emphasis added).

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010

Thus, even if the teachings of Traverstat and Kano are precisely as the Office action suggests (which Appellants do not admit), in view of the network relied upon by the Office action, Vis, Traverstat and Kano still fail to teach or suggest “the persistent memory unit is configured to receive the checkpoint data into a region of the persistent memory unit via a remote direct memory write command from a primary process **through the network interface.**” For this reason alone the rejection should be withdrawn and the claims set for issue.

KANO FAILS TO PROVIDE WHAT VIS AND TRAVERSTAT ARE MISSING

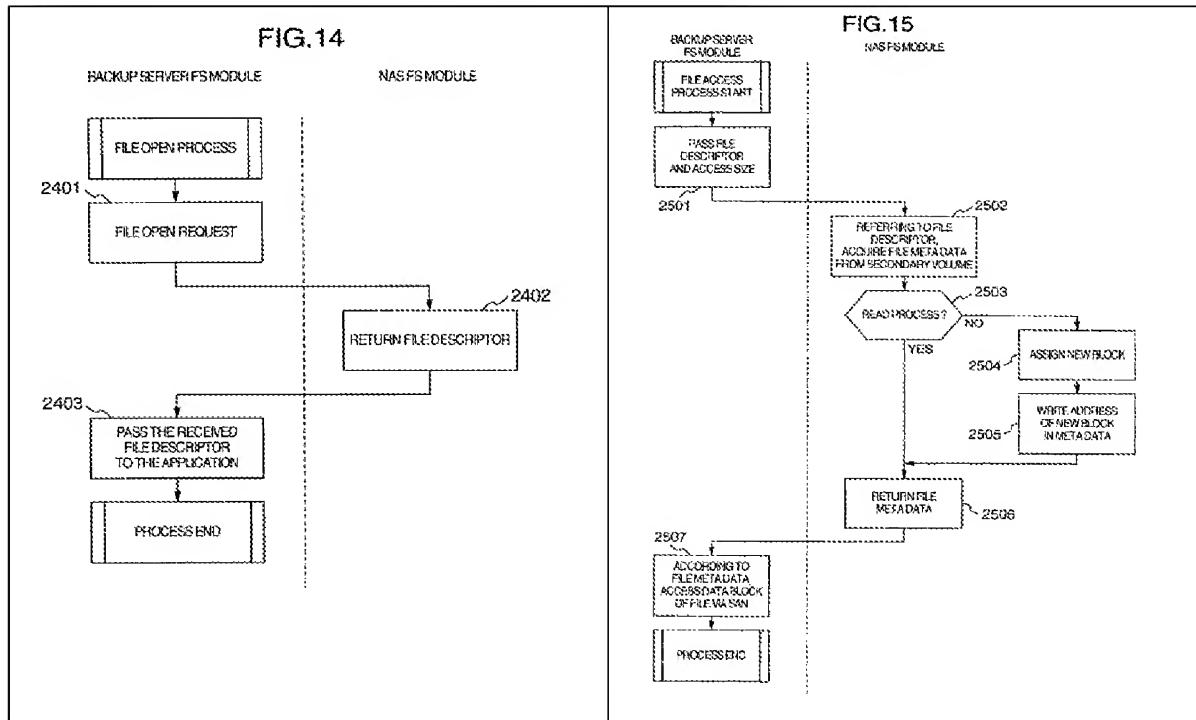
The Office action admits that Vis and Traverstat fall short, and attempts to rely on Kano. In particular, the Office action states:

Kano et al. teach ... a file access for data write the NAS assigns a new data block for storing data (col. 10 *lines* 7-26) and when reception of an open command the backup server issues a file open request to the NAS (col. 9 *lines* 63-67).¹⁰

The portions of Kano relied upon by the Office action deal with file opening procedures, not procedures associated with initializing DMA. Kano’s Figures 14 and 15 are reproduced below for convenience, followed by the sections relied upon by the Office action (along with some additional portions for context).

¹⁰ Office action of September 20, 2010, Page 5, first full paragraph (emphasis original).

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010



The operation of file opening, file access and file list acquisition will be described. In a general file access via a file system, (1) a file is opened to acquire a file descriptor, (2) by referring to the file descriptor, real data of the file is read or an access for write is executed, and (3) after the subject access is completed, the file is closed by referring to the file descriptor. These three processes are executed. In the first embodiment, these three processes are executed by the FS module 145.

...
 A file open process will be described with reference to FIG. 14. Upon reception of an open command, the FS module 250 of the backup server 200 issues a file open request to NAS 100 (2401). If the file exists, the FS module 145 of NAS 100 returns the file descriptor to the backup server 300 (2402). If the file does not exist, the file information may be created from designation made at the time of the file open. If there is no designation, it is regarded as an error. The backup server 300 passes the received descriptor to the application or the like which designated a file open (2403).

Next, a file access process for file read/write will be described with reference to FIG. 15. In order to designate a particular file, the backup server 200 passes the file descriptor and an access size (byte unit) to the FS module 145 of NAS 100 (2501). In accordance

with the file descriptor, the FS module 145 acquires meta data of the file from the secondary volume 130 (2502). As shown in FIG. 18, this meta data 2800 is constituted of an owner 2810, the number 2820 of file hard links, a file size 2830, a data block address LBA 2840, a last access time 2850 and a last update time 2860. If the file access process is for data write (2503), the FS module 145 of NAS 100 assigns a new data block for storing data (2504) and writes LBA of the new block in the meta data 2800 (2505). The FS module 145 of NAS 100 returns the meta data 2800 of the file to the FS module 250 of the backup server 200 (2506). By referring to the address 2840 of the meta data 2800 indicating the file data LBA, the FS module 250 of the backup server 200 accesses the volume via SAN 500 (2507). In the second embodiment, the secondary volume 130 is referred to via SAN 500.¹¹

Even a cursory reading of the cited sections, in view of Figures 14 and 15 reproduced above, reveals that Kano is directed to file level operations on a network attached storage (NAS) device.

Again, representative claim 1 specifically recites:

1. A system for storing checkpoint data comprising:
a network interface to an external network; and
a persistent memory unit coupled to the network interface,
wherein:

the persistent memory unit is configured to receive the checkpoint data into a region of the persistent memory unit via a **remote direct memory write command** from a primary process through the network interface, and to provide access to the checkpoint data in the region via a **remote direct memory read command** from a backup process through the network interface, **wherein the remote direct memory write command is preceded by a create request for the region and the read command is preceded by an open request for the region**; and

the backup process provides recovery capability in the event of a failure of the primary process.

¹¹ Kano Col. 9, line43 through Col. 10, line 26.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

Appellants respectfully submit that Vis, Traverstat and Kano fail to teach or suggest such a system. The Office action attempts to rely on Kano for the teachings regarding the claimed DMA procedures; however, Kano is directed to file level operations, not DMA. The file opening procedures taught by Kano fail to teach, suggest, or even imply a procedure for DMA as claimed.

THE BROADEST REASONABLE INTERPRETATION RULE CANNOT EXPAND THE DMA OPERATIONS CLAIMED TO FILE-LEVEL OPERATIONS

Appellants understand and acknowledge that Examiners are to use the broadest reasonable interpretation of claims when examining for patentability. However, the reasonableness of an interpretation is bounded by Appellants' specification. In particular, the Manual of Patent Examining Procedures (MPEP) admonishes:

During patent examining, the pending claims must be "given their broadest reasonable interpretation **consistent with the specification.**"¹²

Appellants respectfully that the DMA operations claimed cannot be expanded under the guise of the broadest reasonable interpretation to cover the file level operations of Kano. "The protocol of giving claims their broadest reasonable interpretation during examination does not including giving claims a legally incorrect interpretation. This protocol is solely an examination expedient, not a rule of claim construction."¹³

Based on the foregoing, Appellants respectfully submit that the rejections of the claims be reversed, and the claims set for issue.

2. Claim 2

Claim 2 stands rejected as allegedly obvious over Vis, Traverstat and Kano.

¹² MPEP 8th, Rev. 6, August 2007, § 2111, p. 2100-46 (emphasis added).

¹³ *In re Robert Skvorecz*, 580 F.3d 1262 (Fed. Cir. 2009)

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010

Claim 2 recites:

2. The system of Claim 1, further comprising:
a persistent memory manager configured to program the network
interface with information used by the network interface to perform
virtual-to-physical address translation.

Appellants respectfully submit that Vis, Traverstat and Kano fail to teach or suggest the limitations of claim 2. In rejecting claim 2, the Office action states:

Viswanatham et al. does not explicitly teach the system of Claim 1, further comprising: a persistent memory manager configured to program the network interface with information used by the network interface to perform virtual-to-physical address translation.

Traverstat et al. teach the virtual heap with a page table and offset based address translation may be used to convert virtual heap references into in-memory heap references (*col. 19 lines 23-31*).¹⁴

The sole location of Traverstat relied upon for support is reproduced immediately below.

Paging provides a simple model to move data from the persistent store 120 to the in-memory heap 108 in virtual machine 100. In one embodiment, a page table 122 and offset based address translation may be used to convert virtual heap 110 references into in-memory heap 108 references. Relatively small pages may be used to reduce heap waste. In one embodiment, a paging-based approach may enable page protection mechanisms and support for DMA and block I/O devices.¹⁵

However, in even a cursory review of Traverstat, particularly the cited portion of column 19, **Traverstat appears to be silent as to where or what portion of Traverstat performs the translation.** Here again, the Office action is attempt to rely on an inherent teaching of a reference.

¹⁴ Office action of September 20, 2010, Page 5, third and fourth full paragraphs (emphasis original).

¹⁵ Traverstat Col. 19, lines 24-31.

**TRAVERSTAT FAILS TO
INHERENTLY TEACH THE TRANSLATION
PERFORMED BY WAY OF NETWORK INTERFACE**

Again, the Manual of Patent Examining Procedures (MPEP) provides the following guidance as to inherency.

To establish inherency, the extrinsic evidence must make clear that the missing descriptive matter is **necessarily present** in the thing described in the reference, and that it would be recognized by persons of ordinary skill. Inherency, however, may not be established by probabilities or possibilities. **The mere fact that a certain thing may result from a given set of circumstances is not sufficient.**¹⁶

In arriving at the alleged obviousness, the Office action must rely on an allegedly inherent teaching of Traverstat that the heap translation is performed by a network interface. However, the translation of Traverstat could be performed by an array of other devices, such as main processor or the memory controller. Traverstat does not even contain the words “network interface.” It simply does not necessarily present in or from Traverstat that the heap translation taught is could or should performed in a network interface device. Thus, even if the teachings of the Vis and Kano are precisely as the Office action suggests (which Appellants do not admit), Vis, Traverstat and Kano still fail to teach or suggest “a persistent memory manager configured **to program the network interface** with information used by the network interface **to perform virtual-to-physical address translation.**”

Claim 2 is allowable for the same reasons as discussed in Section VII(A)(1) above, but also for the shortcomings discussed in this section.

Based on the foregoing, Appellants respectfully submit that the rejection of the claim 2 be reversed, and the claims set for issue.

¹⁶ MPEP 8th Ed., Rev. 6, September 2007, § 2112(IV), p. 2100-47 (internal quotations omitted, emphasis added).

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

**B. Section 103 Rejection of Claim 9 and 38 Over Vis,
Traverstat, Kano and DeKoning**

Claims 9 and 38 are allowable for the same reasons as discussed in Section VII(A)(2).

C. Conclusion

For the reasons stated above, Appellants respectfully submit that the Examiner erred in rejecting all pending claims. It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's Deposit Account No. 08-2025.

Respectfully submitted,

/mes/

Mark E. Scott
PTO Reg. No. 43,100
CONLEY ROSE, P.C.
(512) 610-3410 (Phone)
(512) 610-3456 (Fax)
ATTORNEY FOR APPELLANTS

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
Legal Dept., M/S 35
3404 E. Harmony Road
Fort Collins, CO 80528

VIII. CLAIMS APPENDIX

1. A system for storing checkpoint data comprising:
 - a network interface to an external network; and
 - a persistent memory unit coupled to the network interface, wherein:

the persistent memory unit is configured to receive the checkpoint data into a region of the persistent memory unit via a remote direct memory write command from a primary process through the network interface, and to provide access to the checkpoint data in the region via a remote direct memory read command from a backup process through the network interface, wherein the remote direct memory write command is preceded by a create request for the region and the read command is preceded by an open request for the region; and

the backup process provides recovery capability in the event of a failure of the primary process.
2. The system of Claim 1, further comprising:
 - a persistent memory manager configured to program the network interface with information used by the network interface to perform virtual-to-physical address translation.

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010

3. The system of Claim 1, wherein the persistent memory unit is configured to provide remote direct memory read access to the checkpoint data to another processor, and the backup process is executed by the other processor.
4. The system of Claim 1, wherein the persistent memory unit provides the checkpoint data through remote direct memory reads by the backup process after the primary process fails.
5. The system of Claim 1, wherein the persistent memory unit is configured to store multiple sets of checkpoint data through remote direct memory writes sent from the processor at successive time intervals.
6. The system of Claim 5, wherein the persistent memory unit provides the multiple sets of checkpoint data through remote direct memory reads upon request by the backup process at one time.
7. The system of Claim 1, wherein the primary process provides the checkpoint data to the persistent memory unit independently from the backup process.
8. The system of Claim 1, wherein the persistent memory unit is configured as part of a remote direct memory access-enabled system area network.

Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010

9. The system of Claim 1, wherein the persistent memory unit is configured with address protection and translation tables to authenticate requests from remote processors, and to provide access information to authenticated remote processors.
38. The system of Claim 1, wherein the persistent memory unit is further configured to store meta-data regarding the contents and layout of memory regions within the persistent memory unit and to keep the meta-data consistent with the checkpoint data stored on the persistent memory unit.
39. The system of Claim 1, wherein the persistent memory unit is further configured to provide access to the checkpoint data in another region via a remote direct memory read command from the backup process through the network interface, wherein the read command is preceded by an open request for the another region.
40. The method of Claim 1, wherein the checkpoint data received by the persistent memory unit overwrites a current set of the checkpoint data.
41. The method of Claim 1, wherein the checkpoint data received by the persistent memory unit is appended to a previous set of the checkpoint data.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

IX. EVIDENCE APPENDIX

None.

**Appl. No. 10/737,374
Second Appeal Brief dated December 6, 2010
Reply to Office action of September 20, 2010**

X. RELATED PROCEEDINGS APPENDIX

None.